

REMARKS

Claims 1-18 are currently pending in the above-identified patent application. Independent claims 1, 9 and 14 have been amended to include the limitations that the throughput of each of said plurality of disk drives is maximized and that each of said plurality of parallel data streams has an approximately equal second throughput. Support for the first of these amendments may be found on page 5, lines 13-15, of the subject Specification, as originally filed, wherein it is stated that: "Because the data streams 112 and 114 may be selected to be the maximum data transfer rate or throughput of the disk drives 118-126, the embodiment 100 may have an overall data throughput of approximately twice that of the individual disk drives 118-126." Support for the second amendment may be found on page 6, line 30, to page 7, line 1, of the Specification, as originally filed, wherein it is stated that: "In a typical multiple disk storage system, the data transfer rates of data streams 320 and 322 will be approximately the same." An obvious typographical error was also corrected in claim 14: "... said first throughput being small than said ~~first~~ second throughput;" No new matter has been added by these changes.

Applicant's attorney, Samuel M. Freund, participated in a Telephonic Interview with Examiner Richard Franklin and Primary Examiner Alford Kindred on November 9, 2007, wherein the subject Office Action was discussed. In particular, Primary Examiner Kindred stated that applicants' argument presented in Amendment A dated July 06, 2007 that independent claims 1, 9 and 14 recite that each of the plurality of parallel data streams is required to have a second throughput equal to that of another data stream is not clear from the language of the claims. Applicant has amended these claims to clarify this issue in response thereto. Further, Examiner Franklin was not persuaded by applicant's argument during the Interview that the Riggles et al. reference teaches away from the present claimed invention by requiring that the bandwidth for each of the disk drives is different, since the Examiner cited Col. 7, lines 16-20, of Riggles et al. as stating that the overall bandwidth of the device of Riggles et al. is the bandwidth of one disk drive times the number of disk drives in the system which teaches the same bandwidth for each parallel data stream. No agreement was reached concerning allowable subject matter.

In the subject Office Action, made final, the Examiner rejected claims 1–18, under 35 U.S.C. 102(b) as being anticipated by U.S. Patent No. 5,724,539 to Riggie et al., hereinafter Riggie, since the Examiner stated that as per claim 1, Riggie teaches a method comprising addressing a plurality of data strips from data to a chosen disk of a plurality of disk drives (Col. 3, lines 6 – 10 and Col. 6, lines 28 – 31); forming a data stream comprising data strips (FIG. 1, Item 90), the data stream having a first throughput (Col. 5, lines 5–8 and 13–17); creating a plurality of parallel data streamers (FIG. 1, Item 110), each of the plurality of parallel data streams having a second throughput (Col. 5, lines 5–8 and 13–17), the second throughput being smaller than the first throughput (Col. 5, lines 5–8 and 13–17); directing the plurality of parallel data streams to a corresponding plurality of the plurality of disk drives (FIG. 1, Item 150 and Col. 6, lines 28–31) such that each data strip in the plurality of data strips is transmitted to the chosen disk of the plurality of disk drives (Col. 6, lines 31–34); and storing each of the data strips on the each of plurality of disk drives (Col.6, lines 31–34).

As per claims 9 and 14, the Examiner continued that Riggie teaches a system comprising a plurality of disk drives (FIG. 1, Item 150) each having a communication channel (FIG. 1, Item 140) capable of communicating at a first throughput (Col. 5, lines 5–8 and 13–17); a controller (FIG. 1, Item 40) adapted to address a plurality of data strips from the data to a chosen disk of the plurality of disk drives (Col. 3, lines 6–10, Col. 6, lines 28–31), and form a data stream comprising the data strips, the data stream having a second throughput (Col. 5, lines 5–8 and 13–17); a buffered switch (FIG. 1, Item 50) in communication with the controller adapted to create a plurality of parallel data streams (Col.6, lines 28–31), each of the plurality of parallel data streams having the second throughput, the first throughput being smaller than the second throughput (Col. 5, lines 5–8 and 13–17); a crossbar switch (FIG. 1, Item 100) in communication with the buffered switch and adapted to direct the plurality of parallel data stream to the plurality of disk drives such that each of the separate data strips are transmitted to each of the plurality of disk drives to which the separate data strips are addressed (Col. 6, lines 31–40); and wherein the plurality of disk drives are adapted to receive the plurality of parallel data streams and store the data strips on the disk drives (Col. 6, lines 31–34).

Dependent claims 2-8, 10-13 and 15-18, were specifically rejected under 35 U.S.C. 102 as being anticipated by Riggle et al.

Applicant respectfully disagrees with the Examiner concerning the rejection of claims 1-18, under 35 U.S.C. 102(b) as being anticipated by U.S. Patent No. 5,724,539, for the reasons to be set forth hereinbelow. Reexamination and reconsideration are requested.

Turning now to Riggle et al., Col. 5, lines 24-41, state: "One of the main objectives of grouping disk drives into an array is to meet the demands for a higher storage subsystem bandwidth. To provide the bandwidth increase in an economically feasible manner the subsystem resources must be used at their optimal capacity levels. If the storage subsystem bandwidth is configured to accommodate the highest transfer rate, the bandwidth capacity is underutilized on average because for full storage capacity utilization data must be placed on all available tracks on the disk surface. Hence a sufficiently large sample transfer unit will span a range of track bands from the disk drives involved. The implication is that the array of disk drives will tend to transfer at an average aggregate bandwidth over a statistically large number of transfers. Having an over-configured subsystem bandwidth is thus undesirable because it results in inefficient and costly resource use of the serial subsystem elements such as the controller buffer shared among the disk drives in the array and the computer host interface bus." (Emphasis added by applicant.).

Clearly, Riggle et al. teaches away from configuring the storage subsystem bandwidth to accommodate the highest transfer rate since according to Riggle et al. the bandwidth capacity is underutilized on average because for full storage capacity utilization, data must be placed on all available tracks on the disk surface.

Column 7, lines 4-43, of Riggle et al. state: "It is possible to format each disk surface 170 of a disk drive 120 to have a constant number of sectors on every track, as shown in FIG. 3. With sector boundaries 180 aligned radially throughout the disk surface 170, each sector 190 is traversed by the read/write head 200 mounted on a positioner 210 and associated with the disk surface in the same period of time. Hence, the same amount of data, such as a 512 byte block, can be written to each track within every sector 190. This uniform sector format of disk surface 170 with the same number

of data bits stored between any two consecutive sector boundaries 180 leads to a constant data transfer rate from any track regardless of its radius. If a transfer unit is distributed among a stripe set of disk drives 130, all the drives can participate in the data movement simultaneously. In the ideal case of a fully parallel transfer the aggregate device bandwidth is thus equal to the individual disk drive bandwidth times the number of drives in the stripe set. However, the uniform sector format of FIG. 3 results in poor utilization of magnetic media storage capacity. Since the length of the track segment bounded by the adjacent sector boundaries 190 increases with track radius, the linear bit density on the outer tracks is lower than that on the shorter inner tracks. An effective way to prevent loss of useful storage capacity is to divide the disk surface into a number of bands of neighboring tracks. As is depicted in a simplified diagram of FIG. 4, each track band is formatted, for example, to have the same number of sectors, each sector defining the length of a track segment needed to store a block of data. An ever greater number of sectors is written on the outer bands in order to maintain a more nearly constant linear bit density throughout the disk surface. As a disk 220 is rotated at a constant angular velocity, however, the linear velocity of the tracks moving past read/write head 200 increases in direct proportion to the track radius. Since the intent of track banding is to produce a uniform linear bit density on every track, the number of data bits encountered by read/write head per unit time grows linearly with its velocity and, therefore, disk radius. Hence, the data transfer rate with read/write head 200 positioned over a track in an inner band 230 is the lowest, increasing for tracks in a middle band 240, and reaching its maximum for transfers involving tracks located in an outer band 250." (Emphasis added by applicants.).

Moreover, claim 1 of Riggle et al. recites in part: "... each track requiring a different track data transfer rate, the tracks being so selected that a sum of the data transfer rates is substantially equal to the bandwidth of the communication channel;"

Thus, the Riggle et al. reference teaches away from equal data transfer rates to each disk drive since the uniform sector format of FIG. 3 thereof results in poor utilization of magnetic media storage capacity.

Subject claim 1, as amended, recites in part: "...addressing a plurality of data strips from said data to a chosen disk of said plurality of disk drives such that the

throughput of each of said plurality of disk drives is maximized; forming a data stream comprising said data strips, said data stream having a first throughput; creating a plurality of parallel data streams, each of said plurality of parallel data streams having an approximately equal second throughput, said second throughput being smaller than said first throughput" (Emphasis added by applicants.). Similar language may be found in independent claims 9 and 14, as amended.

It is known that "Anticipation requires the disclosure in a single prior art reference of each element of the claim under consideration." *Soundsciber Corp. v. U.S.*, 360 F.2d 954, 960, 148 USPQ 298, 301, adopted, 149 USPQ 640 (Ct. Cl. 1966). Subject independent claims 1, 9 and 14, as amended, recite maximizing the throughput of each of said plurality of disk drives. Riggle et al. does not teach this limitation. Rather, Riggle et al. teaches that if one maximizes the throughput of each disk drive, the bandwidth capacity is underutilized on average because for full storage capacity utilization data must be placed on all available tracks on the disk surface. Thus, applicant respectfully believes that Riggle et al. does not anticipate independent claims 1, 9 and 14, as amended.

Further, in *In re Gurley*, 27 F.3d 551, 31 USPQ2d 1130 (Fed. Cir. 1994), it is stated on page 1131 that: "... A reference may be said to teach away when a person of ordinary skill, upon reading the reference, would be discouraged from following the path set out in the reference, or would be led in a direction divergent from the path that was taken by the applicant. The degree of teaching away will of course depend on the particular facts; in general, a reference will teach away if it suggests that the line of development flowing from the reference's disclosure is unlikely to be productive of the result sought by applicant." Article 2141.02 Differences Between Prior Art and Claimed Invention of the Manual Of Patent Examining Procedure, Section VI requires that a prior art reference must be considered in its entirety, i.e., as a whole, including portions that would lead away from the claimed invention. *W.L. Gore & Associates, Inc. v. Garlock, Inc.*, 721 F.2d 1540, 220 USPQ 303 (Fed. Cir. 1983). The teaching of Riggle et al. is clearly that different individual disk drive bandwidth times are required if poor utilization of magnetic media storage capacity is to be avoided.

Therefore, applicant respectfully believes that upon reading the Riggle et al. reference, one having ordinary skill would be discouraged from following the path of selecting equal individual disk drive bandwidth times, and would be led in a direction divergent from the path taken by applicants. Riggle et al. thus teaches away from the present claimed invention.

Dependent claims 2-8, 10-13 and 15-18, were specifically rejected under 35 U.S.C. 102 as being anticipated by Riggle et al. However, since applicant believes that independent claims 1, 9 and 14 from which these dependent claims depend, are patentable over Riggle et al., applicant believes that no further response is required for the dependent claims.

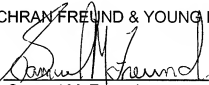
In view of the discussion presented hereinabove, applicant believes that subject claims 1-18, as amended, are in condition for allowance or appeal, the former action by the Examiner at an early date being earnestly solicited.

Reexamination and reconsideration are respectfully requested.

Respectfully submitted,

COCHRAN FREUND & YOUNG LLC

Date: November 13, 2007

By: 
Samuel M. Freund
Reg. No. 30,459
2026 Caribou Drive, Suite 201
Fort Collins, Colorado 80525
Phone: (970) 492-1100
Fax: (970) 492-1101
Customer No.: 27479